

Perils of the One-Size-Fits-All Kernel: A Fast, Secure Search for File System Metadata

LA-UR-20-25810

As we progress beyond the peta-scale storage era, the traditional file metadata management tools and utilities available for distributed and parallel file systems are unable to scale to the billions of files commonly stored in large HPC and cloud data centers. Existing tools that provide parallel capabilities for searching and sifting through large file systems are generally only available for administrators due to the relaxed access controls those tools provide. To better enable users and administrators to locate and manage massive data sets distributed across storage systems and data centers, GUFi (Grand Unified File Index) provides a fast, secure parallel search for locating data sets of interest. GUFi leverages mostly traditional techniques such as embedded databases and indexing to provide a high-performance metadata search capability. However, as we have improved the performance of our metadata indexing service we have identified several bottlenecks within the Linux kernel that limit the performance of our parallel query tools at extreme scales. In this talk we present a detailed description of the GUFi architecture and why it can be securely accessed by users. We then describe why enabling secure user access to the index inevitably runs into performance limitations within current Linux kernels. Finally, we describe an engineering approach for eliminating the kernel-based bottlenecks without requiring additional hardware and describe the performance of GUFi both before and after the application of our performance engineering.